

IN THE CLAIMS

1. (Currently amended) A method of processing a traffic flow in a communication network comprising a plurality of nodes, the method comprising the steps of:

splitting the traffic flow at a given node into a plurality of parts; and

distributing the parts from the given node to respective ones of the plurality of nodes that are designated as participating in a load balancing process for the traffic flow such that each participating node receives a corresponding one of the parts from the given node;

wherein each of at least a subset of the participating nodes receiving one of the parts from the given node routes at least a portion of its received part to one or more destination nodes of the plurality of nodes; and

wherein at least a first one of the participating nodes receiving one of the parts from the given node routes at least a portion of its received part to at least a second one of the participating nodes receiving another one of the parts from the given node.

2. (Currently amended) The method of claim 1 wherein the traffic flow comprises an incoming packet flow arriving at [[a]] the given one of the node[[s]].

3. (Original) The method of claim 1 wherein the traffic flow is split into the plurality of parts in a manner independent of the one or more destination nodes.

4. (Original) The method of claim 1 wherein the traffic flow is split into a plurality of substantially equal parts.

5. (Original) The method of claim 1 wherein the traffic flow is split into a plurality of parts at least two of which comprise non-equal parts.

6. (Currently amended) The method of claim 1 wherein the traffic flow comprises virtually-concatenated data traffic.

7. (Currently amended) The method of claim 1 wherein the traffic flow is split into N parts at a given one of the nodes by maintaining N queues at the given node, and filling the queues from the traffic flow in accordance with a specified queue-filling technique.

8. (Original) The method of claim 7 wherein the specified queue-filling technique comprises one of a round-robin technique and a shortest queue first technique.

9. (Original) The method of claim 1 wherein the traffic flow is split into the plurality of parts utilizing a virtual concatenation technique.

10. (Original) The method of claim 1 wherein the traffic flow is split into the plurality of parts in such a manner that a desired packet format of the traffic flow is maintained in each of the plurality of parts.

11. (Original) The method of claim 1 wherein the parts of the traffic flow are distributed to the respective ones of the participating nodes over pre-provisioned circuits each configured to support a corresponding one of the parts.

12. (Original) The method of claim 1 wherein a given one of the participating nodes routes at least a portion of its received part to a set of destination nodes determined based on destination addresses in packet headers of the portion.

13. (Original) The method of claim 1 wherein if the packet header of a given packet in the part of the flow received by a given one of the participating nodes indicates that the participating node is a final destination node for that packet, the packet is stored in a resequencing buffer of the participating node.

14. (Original) The method of claim 1 wherein if the packet header of a given packet in the part of the flow received by a given one of the participating nodes indicates that the participating node is not a final destination node for that packet, the packet is stored in a particular one of a

plurality of output queues of the participating node that is associated with the final destination node for the packet.

15. (Original) The method of claim 1 wherein at least one of the splitting step and the distributing step is implemented at least in part in software running on a processor of a node or other element of the network.

16. (Currently amended) An apparatus for use in processing a traffic flow in a communication network comprising a plurality of nodes, the apparatus comprising:

a processing device comprising a processor coupled to a memory, the processing device being operative to split the traffic flow at a given node into a plurality of parts, and to distribute the parts from the given node to respective ones of the plurality of nodes that are designated as participating in a load balancing process for the traffic flow such that each participating node receives a corresponding one of the parts from the given node;

wherein each of at least a subset of the participating nodes receiving one of the parts from the given node routes at least a portion of its received part to one or more destination nodes of the plurality of nodes; and

wherein at least a first one of the participating nodes receiving one of the parts from the given node routes at least a portion of its received part to at least a second one of the participating nodes receiving another one of the parts from the given node.

17. (Original) The apparatus of claim 16 wherein the processing device comprises one of the participating nodes of the network.

18. (Original) The apparatus of claim 16 wherein the processing device is implemented as one or more integrated circuits.

19. (Currently amended) An article of manufacture comprising a machine-readable medium storing executable instructions for use in processing a traffic flow in a communication

network comprising a plurality of nodes, the one or more instructions when executed in a processor implementing a method comprising the steps of:

splitting the traffic flow at a given node into a plurality of parts; and

distributing the parts from the given node to respective ones of the plurality of nodes that are designated as participating in a load balancing process for the traffic flow such that each participating node receives a corresponding one of the parts from the given node;

wherein each of at least a subset of the participating nodes receiving one of the parts a given node routes at least a portion of its received part to one or more destination nodes of the plurality of nodes; and

wherein at least a first one of the participating nodes receiving one of the parts from the given node routes at least a portion of its received part to at least a second one of the participating nodes receiving another one of the parts from the given node.